




Search Results

[BROWSE](#)

[SEARCH](#)

[IEEE XPLORE GUIDE](#)

Results for "((store-in <near/4> ((level one or l1 or internal or upper) <near/3> cache) <paragrap..."



Your search matched **0** documents.

A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

» Search Options

[View Session History](#)

Modify Search

[New Search](#)

(((store-in <near/4> ((level one or l1 or internal or upper) <near/3> cache) <paragraph> >>

☐ Check to search only within this results set

» Key

Display Format: ☒ Citation ☐ Citation & Abstract

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

No results were found.

Please edit your search criteria and try again. Refer to the Help pages if you need assistance with your search.



[Help](#) [Contact Us](#) [Privacy & Policy](#)

© Copyright 2005 IEEE - All rights reserved.



Terms used

store in **near/4** **level one** or **L1** or **internal** or **upper near/3** **cache** **paragraph** **duplicate** **tag** and **invalid near/4** **l**

Sort results by

Display results
 [Save results to a Binder](#)

Try ar

 [Search Tips](#)

Try th

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)


Best 200 shown

1 [Parallel execution of prolog programs: a survey](#)


Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo
July 2001

ACM Transactions on Programming Languages and Systems (TOPLAS), Vol

Publisher: ACM Press

Full text available:  [pdf\(1.95 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), :

Since the early days of logic programming, researchers in the field realized the potential for exp of logic programs. Their high-level nature, the presence of nondeterminism, and their referentia make logic programs interesting candidates for obtaining speedups through parallel execution. . applications of logic programming frequently involve irregular computatio ...

Keywords: Automatic parallelization, constraint programming, logic programming, parallelism,

2 [Anatomy of a native XML base management system](#)

T. Fiebig, S. Helmer, C.-C. Kanne, G. Moerkotte, J. Neumann, R. Schiele, T. Westmann

December 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**,

Publisher: Springer-Verlag New York, Inc.

Full text available:  [pdf\(300.97 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index](#)

Several alternatives to manage large XML document collections exist, ranging from file systems specifically tailored XML base management systems. In this paper we give a tour of Natix, a dai scratch for storing and processing XML data. Contrary to the common belief that management (traditional databases like relational systems, we illustrate how almost every component in a ...

Keywords: Database, XML

3 [The V-Way Cache: Demand Based Associativity via Global Replacement](#)



Moinuddin K. Qureshi, David Thompson, Yale N. Patt

May 2005

ACM SIGARCH Computer Architecture News , Proceedings of the 32nd Anr Architecture ISCA '05, Volume 33 Issue 2

Publisher: IEEE Computer Society, ACM Press

Full text available:  [pdf\(231.93 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

As processor speeds increase and memory latency becomes more critical, intelligent design and increasingly important. The efficiency of current set-associative caches is reduced because prog memory accesses across different cache sets. We propose a technique to vary the associativity demands of the program. By increasing the number of tag-store entries relative to the ...

4 [System-level power optimization: techniques and tools](#)



Luca Benini, Giovanni de Micheli

April 2000

ACM Transactions on Design Automation of Electronic Systems (TODAES),

Publisher: ACM Press

Full text available:  [pdf\(385.22 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [...](#)

This tutorial surveys design methods for energy-efficient system-level design. We consider electrical and software layers. We consider the three major constituents of hardware that consume energy: storage units, and we review methods of reducing their energy consumption. We also study models and methods for energy-efficient software design and compilation. This survey ...

5 [Coupling compiler-enabled and conventional memory accessing for energy efficiency](#)




Raksit Ashok, Saurabh Chheda, Csaba Andras Moritz

May 2004

ACM Transactions on Computer Systems (TOCS), Volume 22 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.41 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [...](#)

This article presents Cool-Mem, a family of memory system architectures that integrate conventional aware address translation, and compiler-enabled cache disambiguation techniques, to reduce energy architectures. The solutions provided in this article leverage on interlayer tradeoffs between architectural layers. Cool-Mem achieves power reduction by statically matching memory operations with energy

Keywords: Energy efficiency, translation buffers, virtually addressed caches

6 [Call graph prefetching for database applications](#)



Murali Annavaram, Jignesh M. Patel, Edward S. Davidson

November 2003

ACM Transactions on Computer Systems (TOCS), Volume 21 Issue 4

Publisher: ACM Press

Full text available:  [pdf\(701.71 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [...](#)

With the continuing technological trend of ever cheaper and larger memory, most data sets fit in main memory. In this configuration, the performance bottleneck is likely to be the gap between memory access latency. Previous work has shown that database applications have large instruction processor caches effectively. In this paper, we propose Call Graph Prefetching (CGP), ...

Keywords: Instruction cache prefetching, call graph, database

7 [Processor microarchitecture II: AEGIS: architecture for tamper-evident and tamper-resistant](#)



G. Edward Suh, Dwaine Clarke, Blaise Gassend, Marten van Dijk, Srinivas Devadas

June 2003

Proceedings of the 17th annual international conference on Supercomputing

Publisher: ACM Press

Full text available:  [pdf\(286.90 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [...](#)

We describe the architecture for a single-chip AEGIS processor which can be used to build computer software attacks. Our architecture assumes that all components external to the processor, such as different implementations. In the first case, the core functionality of the operating system is trusted; also describe a variant implementation assuming an untrusted operating system ...

Keywords: certified execution, secure processors, software licensing

8 [External memory algorithms and data structures: dealing with massive data](#)



Jeffrey Scott Vitter

June 2001

ACM Computing Surveys (CSUR), Volume 33 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(828.46 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [...](#)

Data sets in large applications are often too massive to fit completely inside the computers into communication (or I/O) between fast internal memory and slower external memory (such as disk); this article we survey the state of the art in the design and analysis of external memory (or EM); is to exploit locality in order to reduce the I/O costs. We consider a variety of ...

Keywords: B-tree, I/O, batched, block, disk, dynamic, extendible hashing, external memory, I/O methods, multilevel memory, online, out-of-core, secondary storage, sorting

9 Cool-Mem: combining statically speculative memory accessing with selective address tran



Raksit Ashok, Saurabh Chheda, Csaba Andras Moritz

October 2002

**ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , ACM SIG
Proceedings of the 10th international conference on Architectural support
systems ASPLOS-X**, Volume 36 , 37 , 30 Issue 5 , 10 , 5

Publisher: ACM Press

Full text available: [pdf\(1.42 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), :

This paper presents Cool-Mem, a family of memory system architectures that integrate convent aware address translation, and compiler-enabled cache disambiguation techniques, to reduce er architectures. It combines statically speculative cache access modes, a dynamic CAM based Tag mispredicted accesses, various conventional multi-level associative cache organizations, embed

10 Information flow inference for free



François Pottier, Sylvain Conchon

September 2000

**ACM SIGPLAN Notices , Proceedings of the fifth ACM SIGPLAN internation
ICFP '00**, Volume 35 Issue 9

Publisher: ACM Press

Full text available: [pdf\(749.77 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), :

This paper shows how to systematically extend an arbitrary type system with dependency infor interference proofs for the new system may rely upon, rather than duplicate, the soundness prc virtually any of the type systems known today with information flow analysis, while requiring or on an untyped operational semantics for a labelled calculus akin to core ML. Thus, it i ...

11 Memory-wall: Bloom filtering cache misses for accurate data speculation and prefetching



Jih-Kwon Peir, Shih-Chang Lai, Shih-Lien Lu, Jared Stark, Konrad Lai

June 2002

Proceedings of the 16th international conference on Supercomputing

Publisher: ACM Press

Full text available: [pdf\(248.57 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), :

A processor must know a load instruction's latency to schedule the load's dependent instruction processors do not know this latency until well after the dependent instructions should have bee themselves and the load. One solution to this problem is to predict the load's latency, by predic cache. Existing cache hit/miss predictors, however, can only correctly ...

Keywords: bloom filter, data cache, data prefetching, data speculation, instruction scheduling

12 Level set and PDE methods for computer graphics



David Breen, Ron Fedkiw, Ken Museth, Stanley Osher, Guillermo Sapiro, Ross Whitaker

August 2004

Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '0

Publisher: ACM Press

Full text available: [pdf\(17.07 MB\)](#)

Additional Information: [full citation](#), [abstract](#)

Level set methods, an important class of partial differential equation (PDE) methods, define dyr surface) of a sampled, evolving nD function. The course begins with preparatory material that i equations to solve problems in computer graphics, geometric modeling and computer vision. Th several different types of differential equations, e.g. the level set eq ...

13 Fast detection of communication patterns in distributed executions

Thomas Kunz, Michiel F. H. Seuren

November 1997

Proceedings of the 1997 conference of the Centre for Advanced Studies or

Publisher: IBM Press

Full text available: [pdf\(4.21 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), :

Understanding distributed applications is a tedious and difficult task. Visualizations based on pr better understanding of the execution of the application. The visualization tool we use is Poet, a Waterloo. However, these diagrams are often very complex and do not provide the user with th experience, such tools display repeated occurrences of non-trivial commun ...

14

Improving direct-mapped cache performance by the addition of a small fully-associative ca

Publisher: ACM Press

Full text available:  pdf(1.20 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), :

Projections of computer technology forecast processors with peak performance of 1,000 MIPS in 1985 could easily lose half or more of their performance in the memory hierarchy if the hierarchy does not improve. This paper presents hardware techniques to improve the performance of caches. Miss caching policy, cache and its refill path. Misses in the cache ...

15 Fast address lookups using controlled prefix expansion



V. Srinivasan, G. Varghese

February 1999 **ACM Transactions on Computer Systems (TOCS)**, Volume 17 Issue 1

Publisher: ACM Press

Full text available: pdf(258.60 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Internet (IP) address lookup is a major bottleneck in high-performance routers. IP address lookups are performed using longest prefix matching. It is compounded by increasing routing table sizes, increased traffic, and the growth of IPv6 addresses. We describe how IP lookups and updates can be made faster using a set of of techniques including controlled prefix expansion, transf ...

Keywords: Internet address lookup, binary search on levels, controlled prefix expansion, expanded router preformance

16 Let caches decay: reducing leakage energy via exploitation of cache generational behavior



Zhigang Hu, Stefanos Kaxiras, Margaret Martonosi

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Publisher: ACM Press

Full text available: pdf(873.03 KB)

Additional Information: full citation, abstract, references, :

Power dissipation is increasingly important in CPUs ranging from those intended for mobile use, for highend servers. Although the bulk of the power dissipated is dynamic switching power, leakage power is becoming a significant fraction of the total. Chipmakers expect that in future chip generations, leakage's proportion of total chip power will increase. This paper presents several methods for reducing leakage power within the cache memories of the CPU. Be ...

Keywords: Cache memories, cache decay, generational behavior, leakage power

17 Automatic tiling of iterative stencil loops



Zhiyuan Li, Yonghong Song

November 2004 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Vol

Publisher: ACM Press

Full text available: pdf(947.69 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [info](#)

Iterative stencil loops are used in scientific programs to implement relaxation methods for numerical simulations. These loops iteratively modify the same array elements over different time steps, which presents opportunities for improving temporal data locality through loop tiling. This article presents a compiler framework for automatically generating tiled code for iterative stencil loops. The primary objective of improving the cache performance. The article first presents a ...



Keywords: Caches, loop transformations, optimizing compilers

18 Multithreading I: Pointer cache assisted prefetching

Jamison Collins, Suleyman Sair, Brad Calder, Dean M. Tullsen

November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture**

Publisher: IEEE Computer Society Press

Full text available:  pdf(1.21 MB)  Publisher Site

Additional Information: [full citation](#), [abstract](#), [references](#), [full text](#)

Data prefetching effectively reduces the negative effects of long load latencies on the performance of processors. Most processors employ hardware structures to predict future memory addresses based on previous patterns. This paper proposes a new method to employ the actual program code to determine future load addresses for prefetching. This paper proposes the

transitions, to aid prefetching. The pointer cache provides, for a given pointer's ...


- 19 Register file and memory system design: Dynamic addressing memory arrays with physical
Steven Hsu, Shih-Lien Lu, Shih-Chang Lai, Ram Krishnamurthy, Konrad Lai
November 2002 **Proceedings of the 35th annual ACM/IEEE international symposium on Microarchitecture**
Publisher: IEEE Computer Society Press


Full text available:

 pdf(907.32 KB)  Publisher Site

Additional Information: [full citation](#), [abstract](#), [references](#), ...

As pipeline width and depth grow to improve performance, memory arrays in microprocessors increase in physical size, which prolongs the access time due to wiring delay. In order to boost multiple cycles to complete an access. This delays the scheduling of dependent instructions and proposes a different circuit organization to enable fast and slow accesses solely de ...

- 20 Cache decay: exploiting generational behavior to reduce cache leakage power
 Stefanos Kaxiras, Zhigang Hu, Margaret Martonosi
May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual architecture ISCA '01**, Volume 29 Issue 2
Publisher: ACM Press

Full text available:  pdf(1.17 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), ...

Power dissipation is increasingly important in CPUs ranging from those intended for mobile use, for high-end servers. While the bulk of the power dissipated is dynamic switching power, leakage. Chipmakers expect that in future chip generations, leakage's proportion of total chip power will

This paper examines methods for reducing leakage power within the cache memori ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Play](#)